

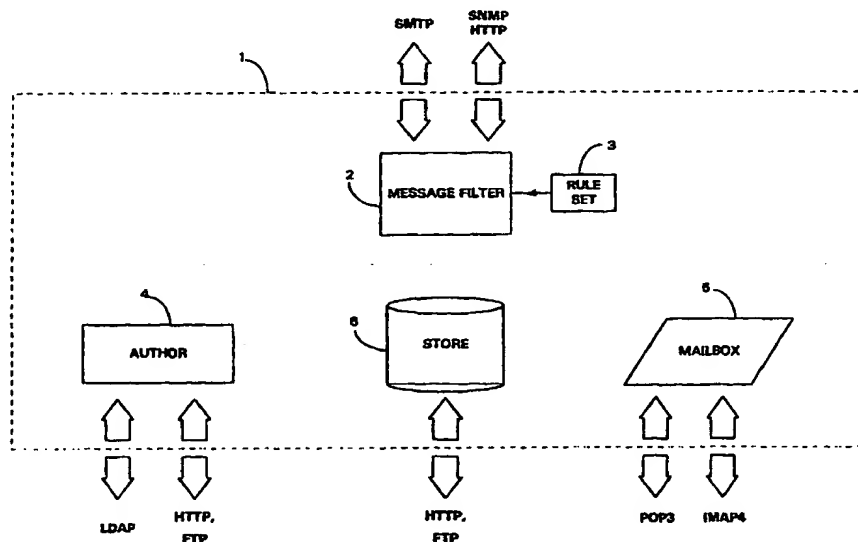


09/909.015

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/58, 29/06		A1	(11) International Publication Number: WO 99/14909
			(43) International Publication Date: 25 March 1999 (25.03.99)
(21) International Application Number: PCT/GB98/02696		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 8 September 1998 (08.09.98)			
(30) Priority Data: 97307200.2 16 September 1997 (16.09.97) EP			
(71) Applicant (for all designated States except US): BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).			
(72) Inventors; and		Published With international search report.	
(75) Inventors/Applicants (for US only): CHALMERS, Jonathan, Andrew [GB/GB]; 53 Verulam Road, St. Albans, Hertfordshire AL3 4DJ (GB). LEGGATT, Jonathan, Stewart [GB/GB]; 12B Fonnereau Road, Ipswich, Suffolk IP1 3JP (GB). SIMS, Neil, James [GB/GB]; 34 Carlford Close, Martlesham Heath, Ipswich, Suffolk IP5 3TB (GB). THYAGARAJAN, Sudha [IN/US]; 14734 Truitt Farm Drive, Centreville, VA 20120-5408 (US).			
(74) Agent: DUTTON, Erica, Lindley, Graham; BT Group Legal Services, Intellectual Property Dept., Holborn Centre, 8th floor, 120 Holborn, London EC1N 2TE (GB).			

(54) Title: MESSAGING SYSTEM



(57) Abstract

A message filter (2) is provided to analyse, according to predetermined criteria of analysis stored in a rules database (3), electronic mail messages being sent or received by an electronic mail system. Depending upon the results of the analysis of a message, the filter (2) may select the whole or a part of the message and arrange for the storage of the selected part in a store (6), constructing a reference message containing a reference to the stored message part and forwarding the reference message to an intended destination of the message.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

MESSAGING SYSTEM

This invention relates to messaging systems and finds particular application in electronic mail systems.

5 A conventional electronic mail system is based typically upon a client/server arrangement. A user of the mail system gains access typically from a personal computer (PC) or workstation running a proprietary mail client product, providing front-end mail facilities to the user. Relatively simple text-based protocols enable the client to communicate with a compatible mail server over an
10 interconnecting network. The mail server provides user access to a personal mailbox in a message store and handles the interchange of messages between users of the mail system and with users of other mail systems. Communication with other mail systems may take place over either a dedicated network, for example where two mail systems are within the same company, or via the Internet
15 or other publicly available network.

A typical mail system includes facilities to enable a user connected to the system to create a message, define the recipient(s) of the message, despatch the message to the defined recipient(s), store incoming messages in a personal mailbox and download received messages from the mailbox to a PC for viewing. Further
20 facilities are typically provided for the management of a user's mailbox.

Most mail system suppliers have adopted standard protocols and message formats for communication between mail system components and with other proprietary mail systems, particularly where the Internet is used to convey messages between mail systems. An accepted standard Internet message header
25 format is described in Internet Request For Comment (RFC) 822 [CROCKER DH, RFC822 - "Standard for the Format of ARPA Internet Text Messages", 1982], published on-line on the Internet, while other protocols such as the Internet Message Access Protocol (IMAP4) [CRISPIN M, RFC1730 - "Internet Message Access Protocol - Version 4", 1994] and Post Office Protocol (POP3) [MYERS J &
30 ROSE M, RFC1725 - "Post Office Protocol - Version 3", 1994] have been adopted for accessing mailboxes.

Mail servers interchange messages with other mail servers of other mail systems, usually via intermediate gateways, using standard message transfer protocols such as the Internet's Simple Message Transport Protocol (SMTP)

described in RFC821 [POSTEL, 1982]. Gateways communicate with other mail system gateways either by direct network interconnection or via a common mail backbone; in the latter case each mail system requires only a gateway to the mail backbone to enable communication with any other mail system with a gateway to the backbone. The Internet's SMTP Mail Service, for example, provides just such a mail backbone.

Conventional electronic mail systems transfer messages by a 'store and forward' method, a message being stored locally before being forwarded to a subsequent stage in a communication path from sender to recipient. Messages are passed from server to gateway, from gateway to gateway and so on using simple protocols such as SMTP. Gateways provide message buffering where necessary to manage the throughput of messages, for example during congestion in an intervening communications link. Gateways also carry out processing where a change in either message format or protocol is required for transmission over the next stage, for instance at a gateway to an X.400 mail backbone or conversion to 7-bit ASCII format at an SMTP gateway, sometimes incurring a considerable processing overhead. If many large messages come together at a particular gateway, considerable delays may arise in the storage and forwarding of those and other messages.

The accepted standard format for messages carried by the Internet mail system is defined by RFC822 and comprises header information, including the identity of the sender and the intended recipient(s), the date and subject, followed by the body of the message in a plain text format. A multi-part message may also be assembled comprising several "body parts", each body part including a header and a body, as defined by RFC822 and Multi-purpose Internet Mail Extensions (MIME) documented in RFC1521 [BORENSTEIN N (Bellcore) & FREED N (Innosoft), RFC1521 - "MIME (Multi-purpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies", September 1993] and RFC1522 [MOORE K, RFC1522 - "MIME (Multi-purpose Internet Mail Extensions) Part Two: Message Header Extensions for non-ASCII Text", University of Tennessee, September 1993], now widely adopted. The MIME standard enables one or more message body parts of various types to be conveyed within the message, each body part including a header and a body encoded as 7-bit ASCII text or other encoding to enable transfer through the message system. The MIME

standard provides a means for specifying, within the respective headers, the type or types of associated message content being conveyed and their method of encoding, e.g. text, bit-mapped images, audio or video samples, within enclosed body parts to the message. The headings enable a receiver of a message to
5 identify the type of data included within a body part and to invoke an appropriate viewer to decode it and reproduce it locally. In this way, multi-media files or a file generated by a particular computer program, a word processor program for example, may be conveyed as part of a conventional mail message within a body part of the message. The body part comprises a header defining the type of data in
10 the enclosed file, and a body comprising the file itself to which has been applied any further encoding necessary for the purposes of transfer through the mail system. A file conveyed within a message body part in this way is known as an "attachment" to the message. For example, a file containing a digitised audio recording may be conveyed as an attachment to a message, the attachment being
15 encoded for transfer by the mail system and being associated with a header defining the type of data in the file.

A typical small organisation or office may use an Internet-connected local messaging system comprising a mail server and a gateway to SMTP mail. Messages are exchanged with other mail systems via the SMTP gateway. Under
20 the store and forward principle, before despatching a message, a local mail server stores it in the sender's local message store. Similarly, messages received from other mail users are stored in the recipient's incoming mailbox in the message store. In a simple arrangement of two directly interconnected mail systems, a message will be stored twice in passing between the sender and a recipient, once
25 by the sender's local mail server and once by the recipient's. A message bound for multiple recipients will be stored by each recipient's local mail server. Further copies of a message may also be stored at intermediate stages in a communication path, such as at gateways, if only on a temporary queuing basis. A message traversing an internal mail system of a moderate sized organisation occupying a
30 number of sites, may pass through three or four store and forward message transfer stages in passing from a sender to a recipient of that message.

While conventional mail systems have the advantage of simplicity of architecture and of the protocols used, such simplicity can be a disadvantage in the absence of unlimited processing, storage and communication bandwidth. The

store and forward principle, while simple, is a relatively inefficient way to handle message transfer, particularly at stages of high demand for message throughput. The potential for total gridlock in a message system is ever present, both within company mail systems and between company and public mail systems. A
 5 malicious user may cause a mail system to become unavailable for days by sending a very large message, such as a 2 GByte video attachment, to many users simultaneously. However, while such "mail bombs" are currently the preserve of the saboteur, there is an increasing demand to be able to send legitimate multimedia message attachments across mail systems. Such attachments are
 10 inherently large, even after data compression, and may not always be reproducible at the intended destination, wasting mail system resources in their delivery.

According to a first aspect of the present invention, there is provided a message filter arranged, in use, to transfer electronic mail messages, the filter comprising:

- 15 means to receive a message to be transferred;
- a rule set defining at least one test to be performed; and
- storing means connected, in use, to a store;

the filter further comprising:

- testing means to perform a test as determined by the rule set;
- 20 selecting means arranged, in dependence upon the result of the test, to select a part of the received message, the storing means being arranged to store that selected part in the store;
- referencing means to assemble a reference message including a reference to the stored part of the message; and
- 25 forwarding means to forward the reference message to an intended destination.

It is known from the MIME standard to construct an electronic mail message containing a reference to one or more files stored in a known file store. The referenced files are not transported through a mail system with the message;
 30 only the reference to the stored file(s) is transported in the message. According to MIME, a mail message may include a header defining the message content type as "message/external-body", and a reference defining the location and identity of an externally stored file. Using this or an equivalent mechanism, a message part may be stored in a file store and a so called "reference message" may be constructed

comprising a reference to the externally stored message part. The process in which a reference to an externally stored message part is conveyed in a mail message, rather than conveying the message part itself, shall be referred to herein as "mail by reference".

5 Preferably, a message filter according to embodiments of the invention is installed in a mail system to selectively implement message transfer using the "mail by reference" process. Thus it is not necessary for a whole message to be transferred to an intended destination by means of the mail system itself - it is only necessary to indicate to a recipient, at each intended destination, where a non-
10 transferred part of the message has been stored. That is, an embodiment of the message filter may intercept a message being sent by the mail system, select a part of the message - a body part conveying an attachment for example, which may be the whole of the message or one of a number of body parts conveying different attachments - extract the selected part, which may include decoding
15 from a particular mail system's encoding if necessary to obtain the selected part, and store the resultant part or attachment in a location accessible by equipment at each intended destination of the message. A conventional message is then assembled - a "reference message" - incorporating a reference to the stored part of the message and sent through the mail system to the intended destination(s) in the
20 normal way, in this example without the stored attachment. In this way an attachment would be said to have been "referenced". On arrival of the reference message at an intended destination, a recipient may access the referenced store and transfer the referenced attachment of the original message from the store using, for example, a file transfer protocol such as FTP, but not using the mail
25 system itself to convey the attachment to its intended destination.

In a preferred embodiment, the at least one test to be performed by the message filter includes identifying the resources available within a mail system to receive messages, determining the resources required to receive the message to be transferred and comparing the resources available with the resources required.

30 Embodiments of the invention may provide a means for analysing messages in a mail system and, subject to various pre-determined conditions, to implement a mail by reference method of message transfer for messages being sent. Messages may be sent by users of the mail system or they may be sent by other pieces of equipment. Mail by reference may be implemented either in respect

of messages bound for destinations within the same mail system or in respect of messages bound for destinations on other mail systems, or both.

Embodiments of the invention may provide a message filter "close to the source" of messages, for example within the mail system serving the originator of large messages or those intended for widespread distribution, behind a mail system
5 "firewall", within a gateway or at some other point, including equipment associated with the originator of messages. This would be of particular benefit within large organisations having several interconnected mail systems serving different sites. Large internal memos may be sent to multiple destinations subject
10 to a pre-determined policy of referencing messages above a certain size, for example. A message filter according to an embodiment of the invention may implement such a policy at a point close to the source of any such memo distribution.

Preferably, the filter may perform a variety of different tests to determine
15 whether a message should be forwarded in its original form or referenced. Pre-determined thresholds may be established for the message filter, for example regarding the total size of a message or the number of intended destinations of a message, taking into account the available storage capacity available to the filter, the message processing capability of the filter and other components of the
20 associated mail system, the number of users of the mail system, etc. If a test on a message to be sent is found to result in a threshold being exceeded in any respect, the filter may reference the whole or a part of the message. Alternatively, the type of data being conveyed by a message may be determined by the filter, the presence of certain types of data causing the whole or a part of the message to be
25 referenced by the filter.

Embodiments of the invention may alternatively provide a message filter "close to the intended destination" of mail messages, for example if such a message filter is installed in each mail system capable of receiving messages, or at a gateway for incoming messages to a mail system or group of mail systems, or
30 installed as part of a firewall to a company mail system. Such a filter may also be installed on equipment associated with particular destinations, for example to assist with the management of incoming messages, or be arranged to act in respect of messages incoming to or outgoing from a particular user or group of users. The message filter may act on messages originating elsewhere in the same

mail system or group of mail systems, on messages incoming to the mail system from other mail systems or on all incoming messages to a particular destination.

According to a second aspect of the invention there is provided a message filter arranged, in use, to receive reference messages, the filter comprising:

5 means to receive a reference message, the reference message including a respective identifier for each of one or more intended destinations and including a reference to a stored message part;

a rule set defining at least one test to be performed; and

storing means connected, in use, to a store;

10 the filter further comprising:

testing means to perform a test as defined by the rule set;

message assembly means arranged, in dependence upon the result of the test, to assemble a message to be sent; and

15 forwarding means to forward the message to be sent to an intended destination of the reference message.

The filter may be arranged to receive reference messages and to deposit a received reference message with an intended recipient without change, or to perform functions in respect of the reference message and the stored part or parts of the original message referenced within the reference message, according to pre-

20 determined conditions established for the filter.

A message filter according to the second aspect of the present invention may be found especially useful, for example, in limiting the amount of storage space consumed within the mail system by large messages bound for many intended destinations within that mail system. Such a facility may, for example,

25 offer protection against malicious attack by "mail bombs" originating either from inside or outside the mail system.

Preferably the store is a distributed data store and the reference to a stored part of a received message comprises an identification of the location of the stored part of the message in the distributed data store.

30 The Internet is a multimedia computer communications network built on worldwide telephone and data networks. Over 100,000 servers of various types are connected to the Internet providing a publicly accessible distributed data store. An "HTTP server" is a particular type of server holding files of information as data written according to an Internet communication protocol called HyperText Transfer

Protocol (HTTP). Data files stored on HTTP servers and accessible by means of HTTP are known as "WEB pages" which form part of the "World Wide Web", or simply the "WEB". WEB pages are written in a special WEB language called HyperText Markup Language (HTML) creating links to other pages on the WEB, as
5 appropriate, and providing a means to navigate through information on the WEB. Information held on the WEB is accessible to anyone having a computer connected to the Internet and with an interest in accessing it. A Universal Resource Locator (URL) has been adopted as a standard among Internet users to provide a consistent international naming convention to identify all WEB resources, including for
10 instance documents, programs, sound and video clips. A URL defines exactly where to find any given resource on the WEB. The Hypertext Transfer Protocol (HTTP) enables URL-identified files (WEB pages) to be accessed and downloaded to user equipment connected to the Internet.

Users typically access information held on the WEB using proprietary WEB
15 browser products running on their personal computers (PCs) or workstations linked to the Internet. WEB browsers communicate with WEB resources, such as HTTP servers, using the standard Internet protocols such as HTTP to download selected WEB pages, interpret embedded HTML commands inserted at the time of markup by the WEB page authors and, if appropriate, display those pages graphically.
20 Browsers are available to reproduce multi-media files transferred over the Internet, offering greater levels of functionality than is typical with electronic mail system client products.

A message filter according to an embodiment of the present invention may store a message or a part of the message on, for example, a commonly accessible
25 WEB server as WEB pages referenced by a URL, constructing and forwarding a corresponding reference message specifying the URL to an intended destination from where, in turn, use may be made of the Internet and Internet protocols to download the URL-referenced message or message part. A recipient at the destination may then use a WEB browser, for example, to view the downloaded
30 message or message part.

Similarly, in respect of an incoming reference message, a message filter according to an embodiment of the present invention may recognise a URL contained in the message and, if required, download the identified WEB page or pages on behalf of a recipient at the destination.

According to a third aspect of the invention there is provided a method of transferring electronic mail messages, comprising:

receiving a message to be transferred, the message including a respective identifier for one or more intended destinations;

5 performing a pre-determined test; and,

in dependence upon the result of said pre-determined test:

selecting a part of the message and storing the selected part in a store;

assembling a reference message, the reference message including a
10 reference to the stored part of the message; and

forwarding the reference message to an intended destination of the received message.

According to a fourth aspect of the invention, there is provided a method of transferring electronic mail messages, comprising:

15 receiving a message to be transferred, the message including a respective identifier for one or more intended destinations and including a reference to a stored message part;

performing a pre-determined test; and,

in dependence upon the result of said pre-determined test:

20 assembling a message to be delivered; and

forwarding the message to be delivered to an intended destination of the received message.

According to a fifth aspect of the invention, there is provided an electronic mail system including a message filter arranged to analyse electronic mail
25 messages, the filter comprising:

means to receive a message to be analysed;

a rule set defining at least one test to be performed; and

storing means connected, in use, to a store;

the filter further comprising:

30 testing means to perform a test as determined by the rule set;

selecting means arranged, in dependence upon the result of the test, to select a part of the received message, the storing means being arranged to store the selected part in the store;

referencing means to assemble a reference message including a reference to the stored part of the message; and

forwarding means to forward the reference message to an intended destination.

5 A message filter will now be described as an embodiment of the present invention, by way of example only, with reference to the accompanying drawings of which:

Figure 1 shows a mail appliance incorporating a message filter according to embodiments of the invention;

10 Figure 2 depicts an example of message transfer by reference;

Figure 3 is a flow diagram showing the operation of a filter according to a particular embodiment of the invention, in respect of outgoing messages from users of a mail appliance incorporating the invention;

15 Figure 4 is a flow diagram showing the operation of a message filter according to a particular embodiment of the invention, in respect of conventional mail messages incoming to users of a mail appliance incorporating the invention;

20 Figure 5 is a flow diagram, extended from Figure 4, showing the operation of a message filter according to a particular embodiment of the invention, in respect of incoming reference messages.

Referring to Figure 1, the message filter 2 is shown as a component of a so called "mail appliance" 1. The mail appliance 1 may be regarded as a logical association of features arranged, according to embodiments of the invention, to provide electronic mail services to a domain of users, such as a company office or
25 to one or more pieces of equipment. The mail appliance 1 may include known features of a conventional electronic mail system, found implemented on a conventional mail server, together with features particular to the present invention. In particular, the mail appliance 1 includes a conventional author component 4 to assist in the creation of electronic mail messages and a conventional mailbox 5
30 providing a store for incoming or outgoing message. It also includes a store 6 for further file storage by the mail appliance 1 and a message filter 2 according to a preferred embodiment of the invention, connected to a rule set 3 to control the operation of the filter 2 for the mail appliance 1. The store 6, while logically part of the mail appliance 1, may be a remote store on a Web server, accessible over the

Internet. The message filter 2 analyses all messages being sent from sources within the domain of the mail appliance 1 or being received by the mail appliance 1 from outside the domain, applying the principles of message transfer and delivery by reference, where required according to the rule set 3.

5 The message filter 2 of Figure 1 may be implemented in one of a number of possible ways. For example, it may be implemented as part of the functionality of a mail server, arranged to intercept messages being sent by clients of the mail server or received by the mail server from elsewhere. Alternatively, the filter 2 may be arranged to operate as a stand-alone device, communicating with mail servers
10 or directly with mail server components using conventional protocols such as SMTP, IMAP4 and POP3 and with remote storage devices using HTTP or FTP for example.

Referring to Figure 2, an example of electronic mail message transfer involving a two part message 20 is shown. The first part 21 is a body part
15 comprising an initial header, "HEADER 1", preferably defined in accordance with the Internet mail standard for header information established by RFC822 in combination with Multi-purpose Internet Mail Extensions (MIME) defined by RFC1521/2, followed optionally by a plain text body inserted by the sender. For example, "HEADER 1" may comprise the following:

20

```
MIME-Version: 1.0
Date:      Wed, 11 Sep 1996 19:08:00 -0400
From:      Aye Sender <asender@this-host.bt.co.uk>
Subject:   Graphic for meeting
25 To:      You There <you.there@that-host.bt.co.uk>
Cc:        The Boss <tboss@some-host.bt.co.uk>
Content-Type: multipart/mixed;
            boundary=99
```

30 The optional plain text portion following "HEADER 1" may be a simple message to the recipient, for example, "Print the attachment and make copies." followed by a delimiter "--99" to define the boundary of the body part.

The second part 22 is a body part comprising a header, "HEADER 2", and a body. The body may be an encoded attachment, for example a document, image

or audio clip being conveyed as part of the message 20. The header "HEADER 2" may comprise simply a description of the attachment data type and method of encoding, for example:

5 Content-Type: image/gif
 Content-Transfer-Encoding: base64

 In this example, the base64 encoded attachment image is inserted following "HEADER 2", terminating with the delimiter "--99--" to indicate the end
10 of the body part and the end of the message 20. If required, additional body parts comprising headers and encoded attachments may be included in the message 20.

 On receiving a message 20 conveying the attachment, the message filter 2, operating in accordance with the rule set 3, may determine that the attachment is, for example, too large to be conveyed by an electronic mail system and should
15 be stored in a store 6, available to the filter 2 and likely to be available from the intended destination of the message. The filter 2 separates the attachment from the message 20, if necessary decoding any encoding applied for the purposes of message transfer, and stores it (24) in the store 6. The filter 2 then inserts or appends to the remaining first part 21 of the message 20 a reference defining the
20 location and identity of the stored attachment (24) in the store 6. The reference may be incorporated as an additional field in the message header or as a simple plain text entry in the message identifiable as such to the recipient. The resultant message 23, including the reference to the stored attachment (24), is called a "reference message". The reference message 23 is then conveyed by the
25 electronic mail system as a conventional message to the intended destination or destinations.

 Preferably, information selected from the message headers "HEADER 1" and/or "HEADER 2" may be stored (24) along with the attachment in the store 6 so that an audit trail may be preserved in respect of stored message parts.
30 Alternatively, the whole of the original message may be stored in the store 6 rather than just an attachment or particular body part of the message.

 Upon receipt of the reference message 23, a recipient user or system at the destination may use the supplied reference in a suitable file transfer mechanism to either transfer or stream the stored attachment (24), via an interlinking

communications network, for local reproduction. The attachment (22, 24) would not, however, be transferred using the electronic mail system.

As illustrated in Figure 1, the filter 2 communicates with mail appliance storage devices (6) and other resources (4, 5) and, if necessary, with other mail systems, using standard network and mail system protocols such as SMTP for message transfer, FTP for file transfer, etc. and supports standard message formats. In addition, a standard mechanism is already available, as defined by MIME, by which a reference to an externally stored message part may be defined and conveyed in a reference message 23 as part of the message header. The MIME "message/external-body" sub-type may be used in the Content-Type field of a MIME message header, followed by a reference to the computing platform, file identity and directory in which the separated message part or attachment is externally stored. Therefore, in the example above, the reference message 23 may comprise a header "HEADER 3" constructed using fields from the original "HEADER 1" and "HEADER 2" together with information on the store location 6 and identity assigned to the stored message part 24, the header being followed by any plain text included with the original message 20 in the first part 21, as follows:

```
20      MIME-Version: 1.0
      Date:      Wed, 11 Sep 1996 19:08:00 -0400
      From:      Aye Sender <asender@this-host.bt.co.uk>
      Subject:    Graphic for meeting
      To:        You There <you.there@that-host.bt.co.uk>
25      Cc:       The Boss <tboss@some-host.bt.co.uk>
      Content-Type: multipart/mixed;
                  boundary = 99
      Content-Type: message/external-body;
                  name = "myslide.gif";
30                  site = "myserver.bt.com";
                  access-type = ANON-FTP;
                  directory = "graphics";
                  mode = "image";
                  expiration = "Mon, 1 May 1997 10:00:00 -0000 (GMT)"
```

in the "graphics" directory as a file with filename "myslide.gif" and may be transferred from the server using the ANON-FTP file transfer protocol. However, other means for specifying the location of a stored message part may be used in a
10 reference message as long as the means are standardised among potential participant mail systems and appliances.

In determining the conditions under which a received message or a part of a message is to be referenced, the filter 2 interrogates the rule set 3. The rule set 3 may comprise a database of rules expressing the policy of the filter in respect of
15 different types of message. Rules may be implemented in a number of different forms dependent upon the complexity of the policy being defined by the rule set and the flexibility required to change the rules from time to time.

For example, a simple policy of referencing any message to be delivered to one of a number of particular mail destinations may be implemented by means of a
20 rule set database comprising a simple plain text data file listing the identities of those destinations in respect of whose messages the policy is to apply. A computer program implementing the functionality of the filter 2 may comprise, at the point at which the rule set 3 is interrogated upon receipt of a message to be delivered, a simple conditional instruction:

25

```
if <MESSAGE_RECIPIENT> in data_file
{
  ## execute MAIL_BY_REFERENCE subroutine if a match has been found
}
else
30 {
  ## execute DEFAULT subroutine if no match has been found
}
```


where MESSAGE_RECIPIENT is the identity of an intended destination specified in the header of the received message and MAIL_BY_REFERENCE is a program subroutine to select a part of the message, store it and create a corresponding reference message. If the identity of a recipient of the message is not contained in the data file, then the filter's DEFAULT program subroutine may pass the message on for delivery without change.

More complex policies may be implemented in the rule set 3 by means of instructions in the rule set database to execute particular filter program subroutines under specified conditions. For example, an instruction may be stored in the rule set 3 to reference all messages addressed to a particular destination you.there@that-server.bt.co.uk, but only if the message size exceeds 10kbytes and there is at least one attachment included in the message. An appropriate instruction may be as follows:

```

15      you.there@that-server.bt.co.uk
        {SIZE:>10k || ATTACHMENTS:>0}MAIL_BY_REFERENCE
        {}DEFAULT

```

The designations "SIZE:" and "ATTACHMENTS:" may refer to predefined functions included in the computer program implementing the functionality of the filter 2, designed respectively to determine the size of the received message, for example by counting the bytes, and to determine the number and type of body parts making up the received message, for example by analysing the message headers. In the event that the conditional statement within brackets { ... } is true, the MAIL_BY_REFERENCE program is to be executed to reference a part or parts of the received message, otherwise to execute a DEFAULT program - to deliver the message intact, for example.

If required, instructions held within a rule set database may become quite complex under this scheme. For example:

```

30      {SIZE:>10k || ATTACHMENTS:=1 || RECIPIENTS:>5};
        MAIL_BY_REFERENCE
        {DESTINATION_DOMAIN:bt.com}MAIL_BY_REFERENCE_USER
        {ATTACHMENTS:>1}MAIL_BY_REFERENCE_ATTACHMENTS

```

{ }DEFAULT

This rule would reference any message of size greater than 10 kbytes, conveying a single attachment and bound for more than 5 destinations.

- 5 Conditionally, messages bound for the domain "bt.com", would only be referenced in respect of particular users, defined elsewhere in the rule set for example. Failing that, if the message included more than one attachment, then each would be referenced in the text portion of the reference message. The DEFAULT behaviour would be carried out by the filter if none of these requirements were met.

- 10 Advantageously, a message filter may be located at an early stage in the transmission path of a message, at the originating mail appliance for example, so that any portion due to be filtered out of an electronic mail system by the filter may be unbound from the message, stored and referenced close to its source rather than be passed through a significant portion of the mail system.

- 15 With reference to Figure 3, a flow diagram is provided showing the operation of a message filter 2 wherein messages are being sent from sources within the domain of a mail appliance 1 incorporating the filter 2. To begin, at STEP 301, the filter 2 receives a message to be sent from within the domain. At STEP 302 the filter 2 interrogates the rule set 3 to determine the tests to be
- 20 performed on the received message, expressed for example using the scheme described above. In the example of Figure 3, the tests defined by the rule set 3 begin at STEP 303 with a determination by the filter 2 of the size of the message to be sent. At STEP 304, the filter 2 compares the message size from STEP 303 with a threshold size, 100 Kbytes for example, as determined by the rule set 3. If
- 25 the size does not exceed the threshold then the filter proceeds to STEP 310 to determine the type of data being conveyed by the message. At STEP 311, the filter 2 tests for a multi-media data type attachment. If the message data type is found not to be multi-media, then the filter proceeds to STEP 320 to determine the number of intended destinations of the message within the domain of the mail
- 30 appliance 1. At STEP 321, the filter 2 compares the number of destinations from STEP 320 with a threshold determined by the rule set 3. If the number is found to exceed the threshold, 20 destinations for example, then at STEP 350 the message is passed on to the next stage of transit; to the incoming mail box (5) associated with each destination if within the domain of the mail appliance 1 or, if external,

sent over the network to the next message transfer stage en route to another mail appliance or mail system.

If, at STEP 304, the message was found to exceed the threshold size for the mail appliance 1 according to the rule set 3, then at STEP 340 the filter 2
5 proceeds to select a part of the message, in this example an attachment to the message, and to store the selected part in a store 6 appropriate to the type of data conveyed by the message. The selected part may, in general, include the whole of the message, or one or more body parts selected from the message. Optionally the selected part or parts include header information. At STEP 341, the filter 2 creates
10 a reference message containing a reference to the location and identity of the stored part of the message in the store 6. At STEP 342 the filter 2 sends the reference message to the or each intended destination of the message. For destinations within the domain of the mail appliance 1, STEP 342 may involve the filter 2 communicating with the mailbox 5 and depositing the reference message in
15 the incoming mail box associated with each nominated destination. In this embodiment, any message part of multi-media data type detected at STEP 311 is also stored and referenced by steps 340 to 342, though a video server, for example, may then be selected within the store 6 at STEP 340 if the data type was found to be 'video'.

20 If the number of intended destinations within the domain of the appliance 1 is found to exceed the mail appliance threshold, at STEP 321, according to the rule set 3, then at STEP 330 the filter 2 determines the total message storage space required to deliver the message to the incoming mail box associated with each nominated destination within the domain. At STEP 331 the filter 2 compares
25 the total storage space required with a threshold level for the mail appliance as determined by the rule set 3. The threshold may be an absolute limit, 3 Mbytes for example, or it may be set by a determination of available mailbox 5 storage space. If, at STEP 331, the threshold is exceeded, then at STEP 340 a part of the message is selected and extracted and a single copy stored in a store 6. STEP 341
30 creates a reference message in respect of that stored part and STEP 342 sends the reference message to each intended destination. It is also possible, if required at STEP 342 according to the rule set 3, to send reference messages only to intended destinations within the domain of the mail appliance 1, the original message being sent in full to destinations outside the domain without referencing.

message throughput may be adapted to rapid operation in a fast processing environment whereas a filter adapted to run on a lap-top PC in support of a single roving mail user may be satisfactorily implemented on a much smaller scale, 5 although offering similar features. Correspondingly, the associated rule set 3 may be adapted to the filter's environment, establishing thresholds for message analysis taking account, for example, of the volume of storage space available to the mail appliance 1, the processing capability of the filter, the number of users within the domain, the expected number of messages incoming to or outgoing from the 10 domain and the communications bandwidth available to and from the mail appliance 1. Available storage space may be automatically monitored so that thresholds relating to message size and number of recipients, for example, may be adjusted dynamically.

With reference to Figure 4, a flow diagram is provided showing the 15 operation of a message filter 2 in respect of messages being received by the mail appliance 1 from outside the domain of the appliance, each message being bound for a mail destination within the domain. At STEP 401 an incoming message is received by the filter 2. At STEP 402 the filter 2 interrogates the rule set 3 to determine the tests to be performed in respect of messages incoming to the mail 20 appliance 1. Specifically in this example, the tests so determined begin at STEP 403 with a test to establish whether the message being received by the appliance 1 and hence by the filter 2 is of the conventional type, that is, all parts of the incoming message are being conveyed and delivered to the appliance 1 without referencing, or whether the message is a reference message containing a reference 25 to a part of the message stored elsewhere and not conveyed to the appliance 1. If at STEP 403 the message is found to be of the conventional type, then the filter determines, at STEP 410, the size of the message. At STEP 411 the filter 2 compares the size from STEP 410 with a threshold size, 100 Kbytes for example, as determined by the rule set 3. If the size is below the threshold then the filter 30 proceeds to determine, at STEP 420, the number of intended destinations of the message located within the domain of the mail appliance 1. At STEP 421, the filter 2 compares the number of destinations from STEP 420 with a threshold number, 20 for example, as determined by the rule set 3 for the appliance 1. If the number is below the threshold then at STEP 450 the filter communicates with the mailbox

5 of the appliance 1 and deposits the message in the incoming mail box associated with each nominated destination.

If, at STEP 411, the message size is found to exceed the threshold, then at STEP 440 the filter extracts a part of the message and stores it in an appropriate 'remote' storage device. The size threshold used at STEP 411 may, for example, indicate the amount of 'local' storage available within the store 6 of the appliance 1 which, if exceeded by the message size, or if a threat is perceived of a number of messages arriving of such a size, then 'remote' storage is selected at STEP 440 to protect the appliance against actual or potential overload. A 'remote' component of the store 6 may be provided for such a purpose. At STEP 441 a reference message is created in respect of the stored part of the message, containing a reference to the location of the stored part, and the reference message is delivered, at STEP 450, to the incoming mail box associated with each nominated destination.

If, at STEP 421, the filter determines that the number of intended destinations within the domain is in excess of the threshold for the appliance 1, for example, then a determination is made, at STEP 430, of the total storage space required to deliver the message to each of the nominated destinations, calculated in this example by multiplying the number of nominated destinations within the domain from STEP 420 by the size of the message from STEP 410. At STEP 431 the filter 2 compares the space calculated at STEP 430 with a threshold as determined by the rule set 3, an absolute value of 3 Mbytes for example, or the available storage space in the mailbox 5. If the total size exceeds the threshold, then at STEP 440 a part of the message is extracted, as before, and stored 'remotely'. A reference message is created at STEP 441 and delivered at STEP 450 to the incoming mail box associated with each nominated destination within the domain. Below the threshold for total storage, at STEP 431, the message is delivered intact to each nominated destination (STEP 450).

If, at STEP 403 in Figure 4, the filter 2 finds that a reference message has been received by the mail appliance 1, then further processing steps are carried out by the filter 2 according to the policy expressed by the rule set 3 in respect of reference messages. Preferably, in this embodiment of the invention, the steps are demonstrated by the flow diagram of Figure 5.

Referring to Figure 5, the filter 2 firstly determines, at STEP 501, the location of the stored part or parts of the message from the reference or references contained in the reference message, and the size of those stored parts. If necessary, within STEP 501, the filter 2 may send a signal to the referenced
5 store to obtain information about the size of stored parts. At STEP 502, if it is possible for the filter 2 to determine from the reference that the message part is stored on an overseas server, then, in the interests of economy of cost of access and potential time to transfer the message part, for example, the filter 2 may arrange, at STEP 570, to transfer the stored message part to a more local store,
10 for example the store 6 of the mail appliance 1. This policy may prove particularly advantageous if a number of destinations within the domain are nominated to receive the message and each is likely to require access to the stored message part.

Having moved the stored message part to a more local store at STEP 570,
15 the filter 2 proceeds at STEP 571 to amend the reference message to reference the newly stored part. The same policy, as enacted by steps 570 and 571, may be applied if the filter 2 determines at STEP 510 and 511 that, whereas at STEP 502 the store was not overseas, the store is nevertheless inaccessible from one or more of the nominated message destinations in the domain of the mail appliance 1.
20 In that event, the filter 2 may arrange, at steps 570 and 571 respectively, to transfer the stored part to a store that is accessible from all the nominated destinations, the store 6 of the mail appliance 1 for example, and amend the reference message accordingly.

The next step, in this embodiment, STEP 520, is to determine whether or
25 not the total size of the message, including stored parts, exceeds a threshold size, 100 Kbytes for example, as determined by the rule set 3. If STEP 520 finds that the threshold size is exceeded, then at STEP 580 the filter 2 communicates with the mailbox 5 of the appliance 1 and deposits only the reference message in the incoming mail box associated with each nominated destination. If at STEP 520 the
30 message size does not exceed the threshold, then at STEP 530 the filter 2 proceeds to determine the number of intended destinations of the message within the domain of the mail appliance 1. At STEP 531 the filter 2 compares the number determined from STEP 530 with a threshold, 20 for example, as determined by the rule set 3 for the appliance 1. If the threshold is exceeded then at STEP 580 the

filter 2 communicates with the mailbox 5 of the appliance 1 and deposits only the reference message in the incoming mail box associated with each nominated destination.

If, at STEP 531, the number of destinations is found to be below the
5 threshold, then at STEP 540 the filter 2 proceeds to determine the total storage space required in the mailbox 5 to deliver the full message in conventional format - that is, the stored part or parts included as a part of each copy of the message to be delivered - to each of the nominated destinations in the domain. At STEP 541 the filter 2 compares the total storage space required with a threshold, an absolute
10 value of 3 Mbytes for example, as determined by the rule set 3 for the mail appliance 1, or the actual storage capacity available in the mailbox 5 at the time. If the required space is found to exceed the available storage capacity in the mailbox 5, for example, then at STEP 580 only the reference message is delivered to the incoming mail box associated with each nominated destination. If the required
15 space is found at STEP 541 to be below the threshold, and if sufficient storage capacity is available, the option is available, at STEP 550, for a nominated destination to receive the message in either a referenced format, or in full as a conventional message, the stored part or parts being included as part of the message to be delivered to that destination. STEP 550 determines the preference
20 of each nominated recipient for conventional or reference-type messages. A preference for reference messages at STEP 550 results, at STEP 580 in only the reference message being delivered to the incoming mail box associated with that destination. However, to satisfy a preference at STEP 550 by a particular destination for conventional message formats, the filter 2 first obtains, at STEP
25 560, the stored part or parts of the message from the referenced store. At STEP 561 the filter 2 then reassembles a message comprising any body parts of the reference message along with the stored part or parts retrieved at STEP 560 from the store, encoded to the applicable message standard and including an appropriate header or headers. At STEP 580 the full message is then deposited in the incoming
30 mail box associated with that destination.

A message filter according to particular embodiments of the invention may be introduced into any one of a variety of different mail system types, be they implemented in the form of a mail appliance as described above, or using conventional mail servers. A conventional mail server or a mail appliance may be

adapted to include a message filter according to preferred embodiments of the invention by means of a controlling computer program, written, for example, in the 'C' or "C++" programming languages, and installed onto a controlling computer within the mail system, the mail server for example. Alternatively, in the case of a mail appliance, the functionality provided by the filter may be logically contained within the entity that is the mail appliance, although the controlling software may be physically distributed across one or more computing platforms. In either implementation, the controlling program may be arranged to interface with component parts of a mail system or mail appliance by generating standard protocols such as SMTP, FTP, HTTP, SNMP and IMAP4, appropriate to the component being addressed by the control program. In this way a message filter according to preferred embodiments of the invention may communicate with, for example, a message store to deposit a reference message bound for a particular destination, or a more general store for storing or retrieving attachments.

Preferably, a message filter is linked to a store comprising a Lotus Notes server ("Lotus Notes" is a trade mark of Lotus Development Corporation) on which a Lotus Notes mailbox and a document database is provided. The Lotus Notes server is chosen for its accessibility from destinations of messages being processed by the filter. Lotus Notes allows groups of users to interact and share information that can be of a highly unstructured nature. In particular, it provides a document database server which stores, and manages multi-user access to semi-structured data, including text, images, audio and video. The filter may store, in the Lotus Notes mailbox, a copy of a received message which, as a result of tests performed by the filter, is to be transferred "by reference". Controlling software installed on the Lotus Notes server, optionally implemented within Lotus Notes itself and forming part of the functionality of the filter, then analyses the stored message, selects and copies that part or those parts to be referenced, which may include the whole of the stored message, and creates WEB pages in the document database provided on the Lotus Notes server or to be stored elsewhere, incorporating the selected part or parts and/or including appropriate HTML links to them within the pages. The controlling software on the Notes server then signals a message within the filter indicating the URL(s) of the created WEB pages together with information extracted from the message headers required to identify the intended destinations of the message and any other information, as required. From

this information the filter is able to create a reference message including the URL(s) and to forward it to each of the one or more intended destinations, either by way of an individual copy to each destination or, if bound for more than one destination, alternatively as a single copy addressed to multiple destinations.

- 5 In the event that a reference message is lost or, for some reason, does not reach an intended destination, the filter may create another reference message from the information held within the Lotus Notes mailbox on the Lotus Notes server and attempt to sent it.

- Any thresholds or other filtering actions defined for the filter 2 according
10 to the rule set 3 may be arranged to apply only in respect of a single user or group of users within the domain of the mail appliance 1 or mail system, or in respect of a user or group of users of another mail appliance or mail system. The rule set 3 may set different thresholds or filtering actions in respect of messages to or from different users or groups of users. If the mail appliance 1 serves a domain
15 comprising a single user, then the rule set 3 will, naturally, define actions of the filter 3 in respect of messages outgoing or incoming to that user. However, the rule set 3 in such a case may also take account of specific sources or destinations outside of the domain of the appliance 1 in respect of messages incoming to or outgoing from the single user in the domain of the appliance 1.

- 20 Preferably in embodiments of the filter, to prevent the accumulation of old message parts stored in a store 6, a date of expiration may be specified by the filter 2 and communicated to the store 6 at the time of storing the message part according to a pre-determined policy defined by the rule set 3 for the filter 2. Under store management procedures, the date of expiration of a stored message
25 part is monitored and, when the expiration date arrives, the stored part is automatically archived and/or deleted from the store. The date of expiration may be also be included by the filter 2 in the corresponding reference message forwarded to each intended destination so that a recipient at each destination is informed of the date before which the stored message part is to be read or copied
30 if it is not to become unavailable.

Other possible embodiments of a filter of the present invention may be implemented involving further means of delivery of reference messages to intended destinations and recipients. For example, the filter may be arranged to assemble a synthesised voice reference message indicating the location of a separately stored

part of a message, initiating a process of delivery of the voice reference message by means, for example, of a conventional fixed or mobile telephone call or voicemail message. The reference may alternatively be assembled and delivered as a facsimile message, or as a radio pager message to a pager with an alphanumeric display or voice delivery means, or by any other means.

A filter according to embodiments of the invention may also be involved as a stepping stone in passing a message from one mail system to another mail system, for example as the first point of receipt in a network of internal company mail systems for messages received from outside the company. If required, such a filter may be configured to recognise an intended destination mail system to which a message received by the filter is to be sent, and to apply a policy defined by an appropriate rule set to such messages based upon information held or obtained relating to the identified target mail system or as determined by overall company policy.

CLAIMS

1. A message filter arranged, in use, to transfer electronic mail messages, the filter comprising:
- 5 means to receive a message to be transferred;
a rule set defining at least one test to be performed; and
storing means connected, in use, to a store;
- the filter further comprising:
- testing means to perform a test as determined by the rule set;
- 10 selecting means arranged, in dependence upon the result of the test, to select a part of the received message, the storing means being arranged to store that selected part in the store;
- referencing means to assemble a reference message including a reference to the stored part of the message; and
- 15 forwarding means to forward the reference message to an intended destination.
2. A message filter as in Claim 1, wherein the at least one test to be performed includes identifying the resources available within a mail system to
- 20 receive messages, determining the resources required to receive the message to be transferred and comparing the resources available with the resources required.
3. A message filter as in Claim 1 or Claim 2, wherein the at least one test to be performed includes a test to determine the size of a part of the received
- 25 message and to compare the determined size with a pre-determined threshold.
4. A message filter as in any one of Claims 1 to 3, wherein the at least one test to be performed includes a test to determine the number of intended destinations of the received message and to compare the determined number with
- 30 a pre-determined threshold.
5. A message filter as in any one of Claims 1 to 4, wherein the at least one test to be performed includes a test to determine whether or not the data conveyed by a part of the received message is of a pre-determined type.

6. A message filter as in any one of the preceding claims, wherein the part of the message to be stored is an attachment of the message.
- 5 7. A message filter as in any one of the preceding claims, wherein the part of the message to be stored includes information contained in a header of the message.
8. A message filter as in any one of the preceding claims, wherein the part of
10 the message to be stored is the whole of the received message.
9. A message filter as in any one of the preceding claims, wherein the store is a distributed data store and the reference to a stored part of the received message comprises an identification of the location of the stored part of the
15 message in the distributed data store.
10. A message filter as in any one of the preceding claims, wherein the store is located on a server connected to the Internet and the reference to a stored part of the received message comprises the Internet Universal Resource Locator (URL)
20 of the stored part of the message.
11. A message filter arranged, in use, to receive reference messages, the filter comprising:
- means to receive a reference message, the reference message including a
25 respective identifier for each of one or more intended destinations and including a reference to a stored message part;
- a rule set defining at least one test to be performed; and
- storing means connected, in use, to a store;
- the filter further comprising:
- 30 testing means to perform a test as defined by the rule set;
- message assembly means arranged, in dependence upon the result of the test, to assemble a message to be sent; and
- forwarding means to forward the message to be sent to an intended destination of the reference message.

12. A message filter as in Claim 11, wherein the assembly means are arranged to retrieve the message part referenced by the received reference message and, in dependence upon the result of the test, to assemble a message to be sent including the message part.

5

13. A message filter as in Claim 11, wherein the assembly means are arranged to retrieve the message part referenced by the received reference message, said storing means are arranged to store the message part in the store and the assembly means are further arranged to assemble a message to be sent including a
10 reference to the retrieved message part stored in the store.

14. A method of transferring electronic mail messages, comprising:
receiving a message to be transferred, the message including a respective identifier for one or more intended destinations;
15 performing a pre-determined test; and,
in dependence upon the result of said pre-determined test:
selecting a part of the message and storing the selected part in a store;
assembling a reference message, the reference message including a
20 reference to the stored part of the message; and
forwarding the reference message to an intended destination of the received message.

15. A method of transferring electronic mail messages, comprising:
25 receiving a message to be transferred, the message including a respective identifier for one or more intended destinations and including a reference to a stored message part;
performing a pre-determined test; and,
in dependence upon the result of said pre-determined test:
30 assembling a message to be delivered; and
forwarding the message to be delivered to an intended destination of the received message.

16. A method of transferring electronic mail messages as in Claim 15, including, in dependence upon the result of said pre-determined test, retrieving the stored message part referenced by the received message and including the retrieved message part in the message to be delivered.

5

17. A method of transferring electronic mail messages as in Claim 15, including, in dependence upon the result of said pre-determined test:

retrieving the stored message part referenced by the received message;

storing the retrieved message part in a store; and

10 including a reference to the stored retrieved message part in the message to be delivered.

18. An electronic mail system including a message filter arranged to analyse electronic mail messages, the filter comprising:

15 means to receive a message to be analysed;

a rule set defining at least one test to be performed; and

storing means connected, in use, to a store;

the filter further comprising:

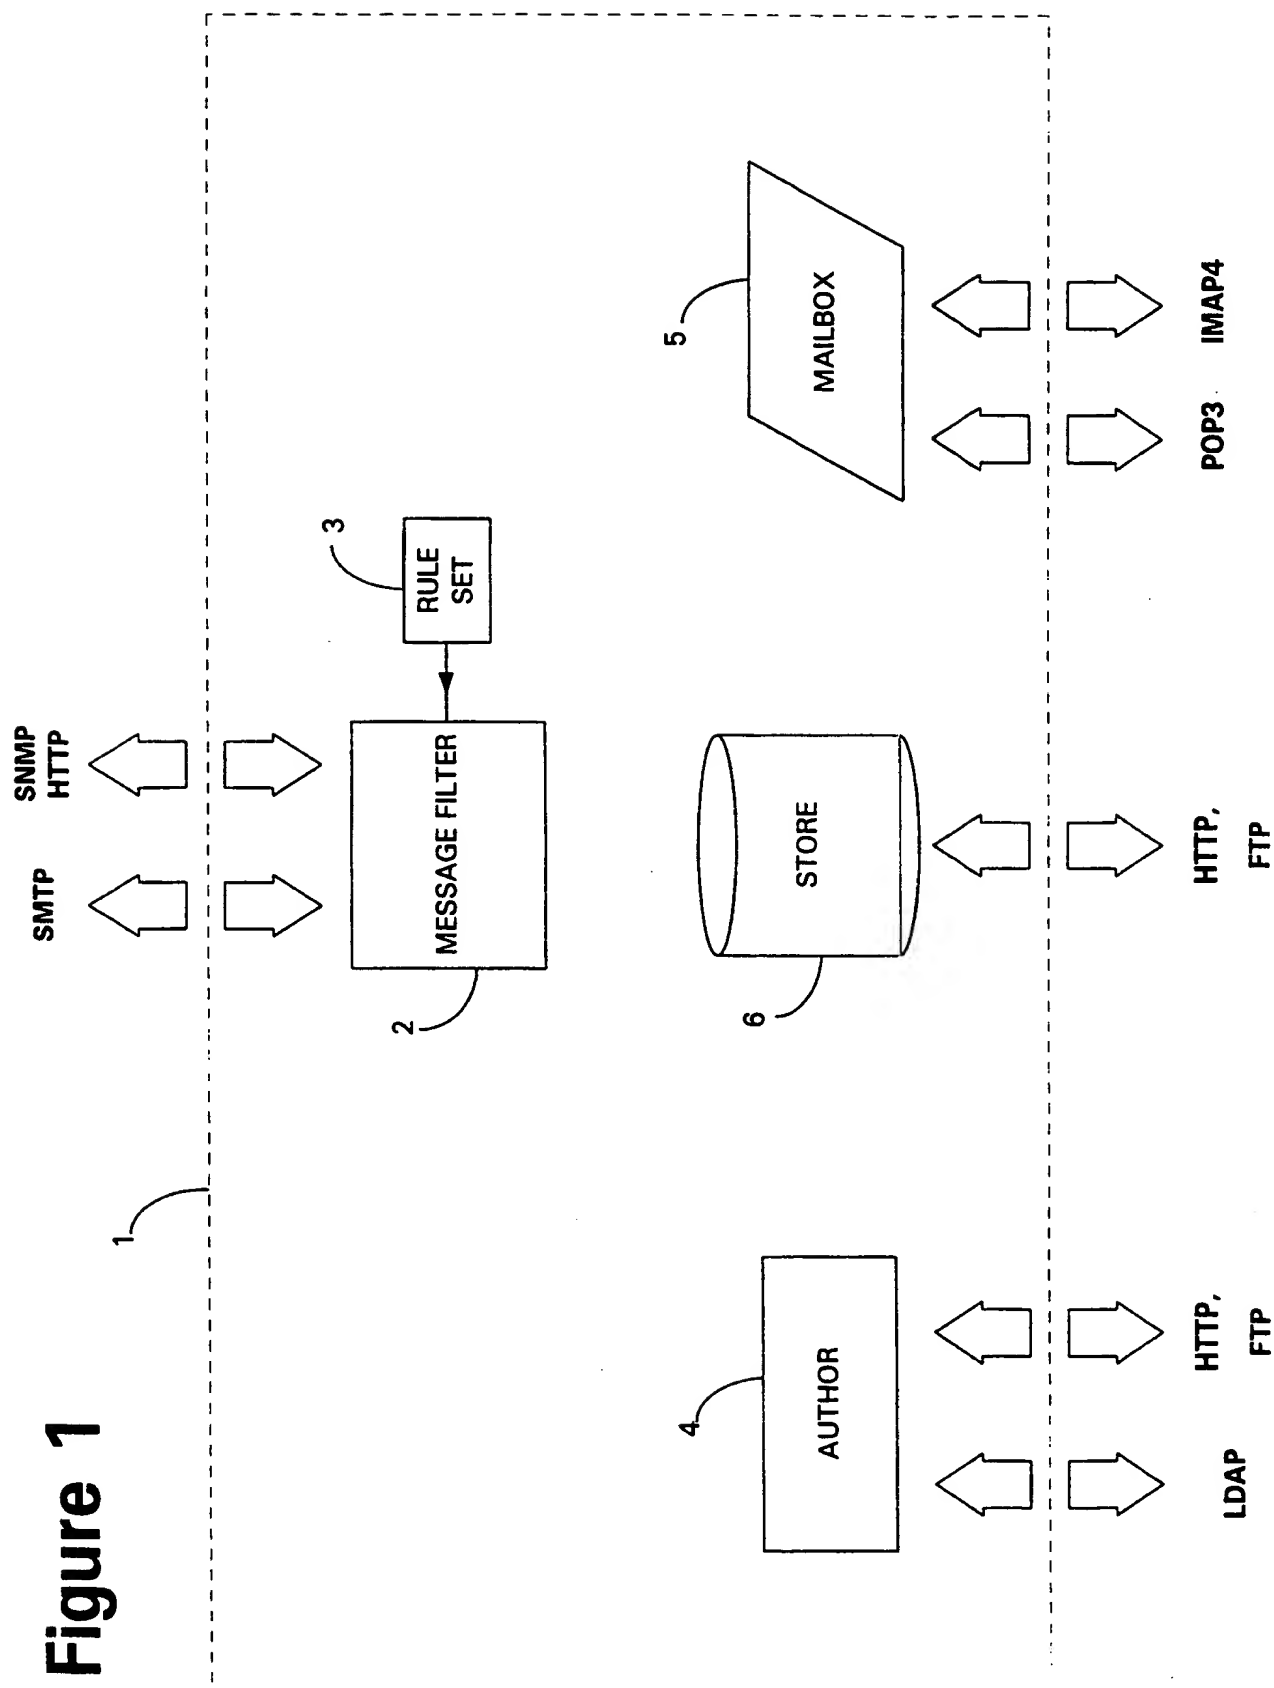
testing means to perform a test as determined by the rule set;

20 selecting means arranged, in dependence upon the result of the test, to select a part of the received message, the storing means being arranged to store the selected part in the store;

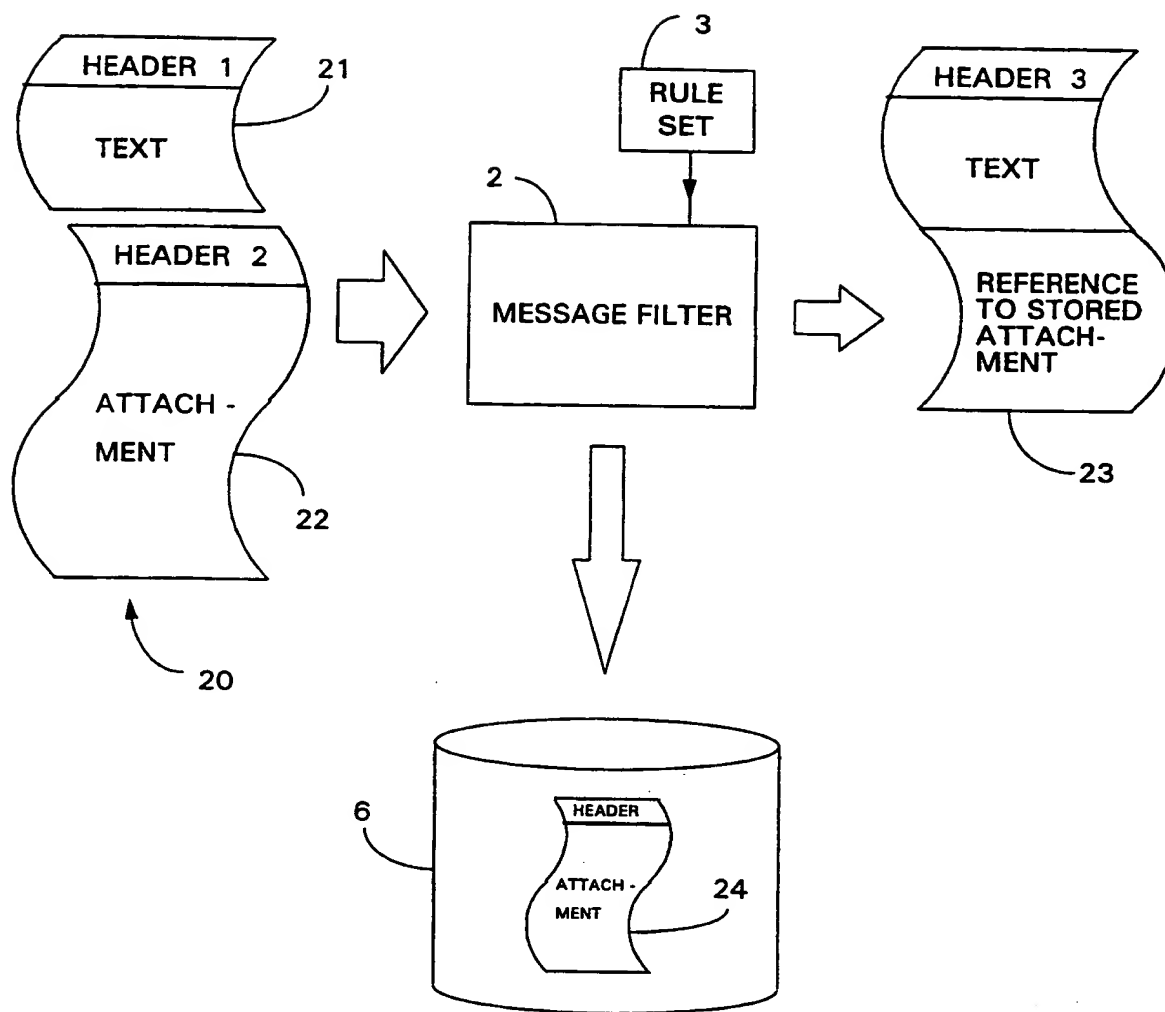
referencing means to assemble a reference message including a reference to the stored part of the message; and

25 forwarding means to forward the reference message to an intended destination.

1/5

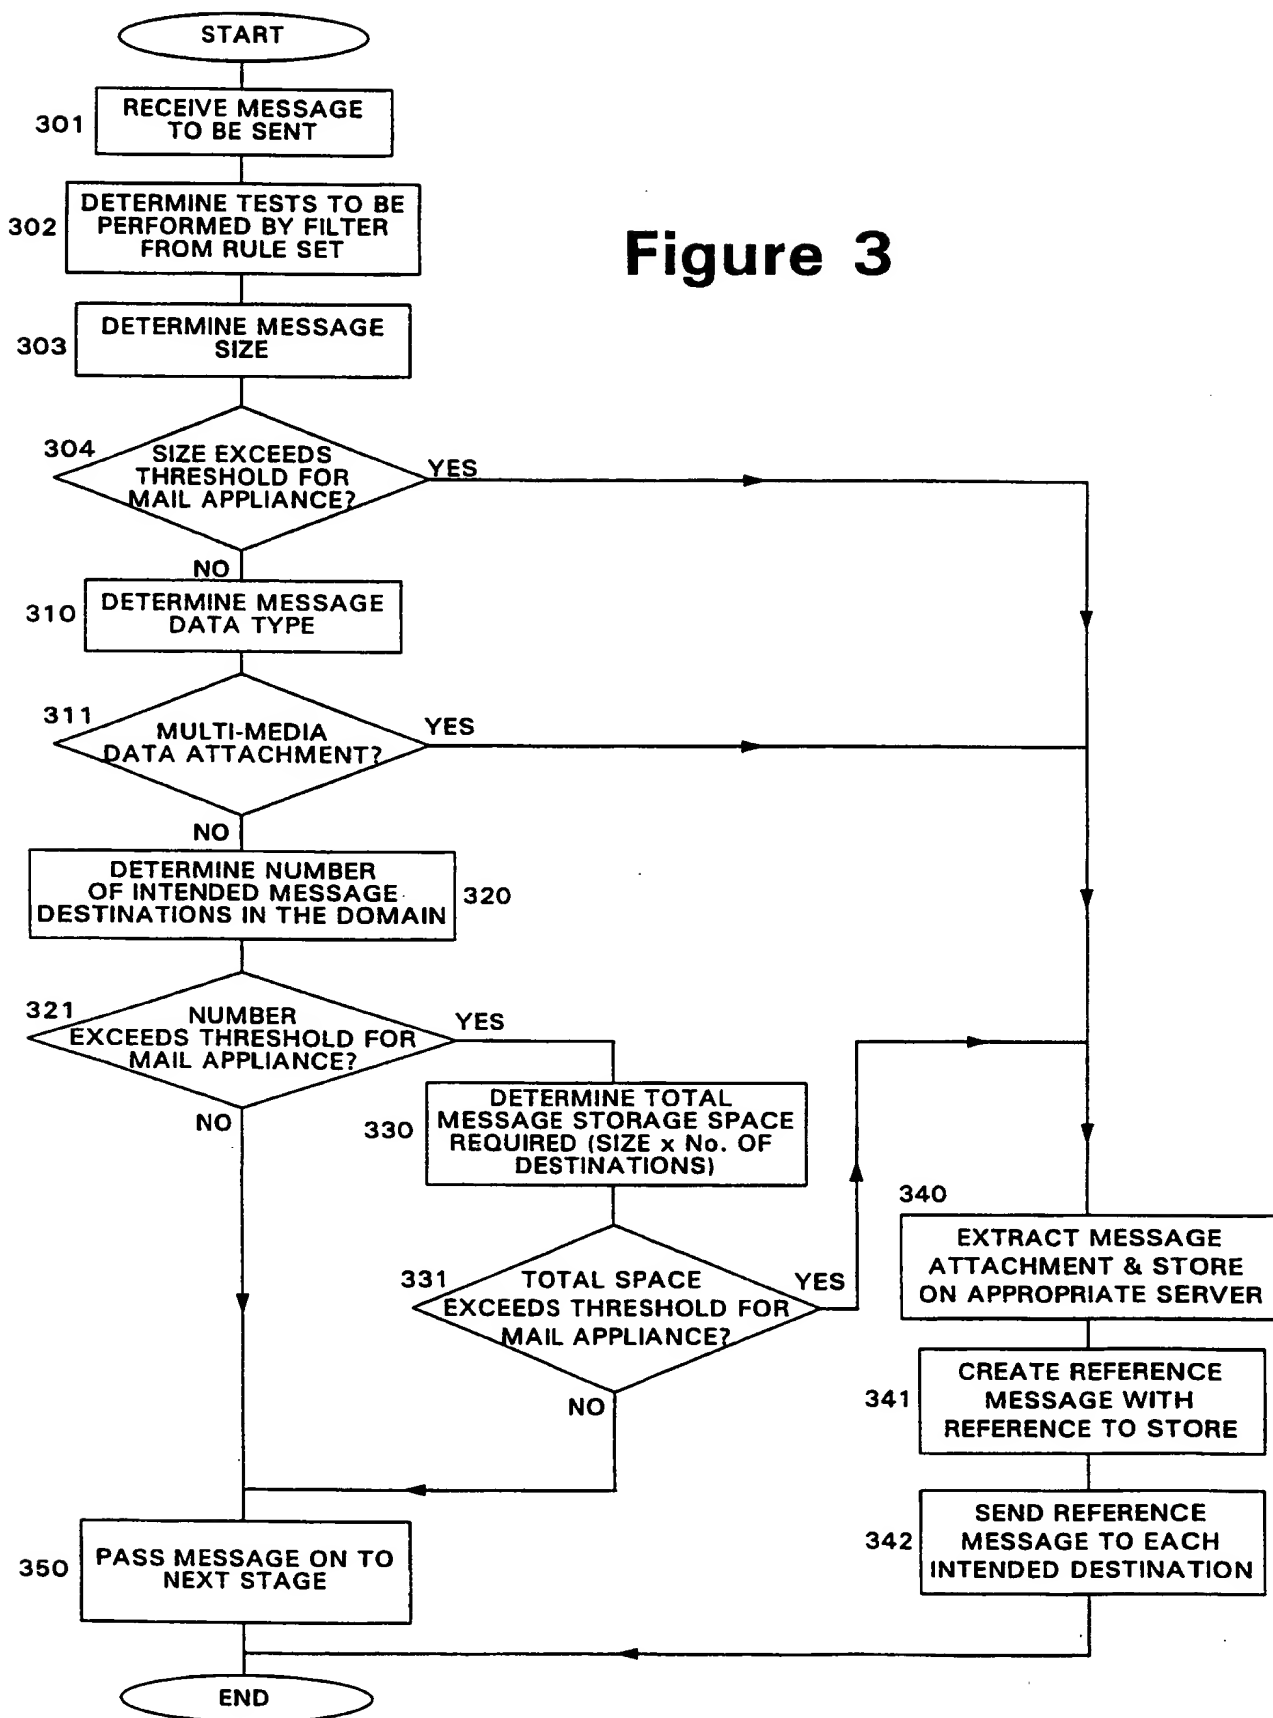


2/5

**Figure 2**

3/5

Figure 3



4/5

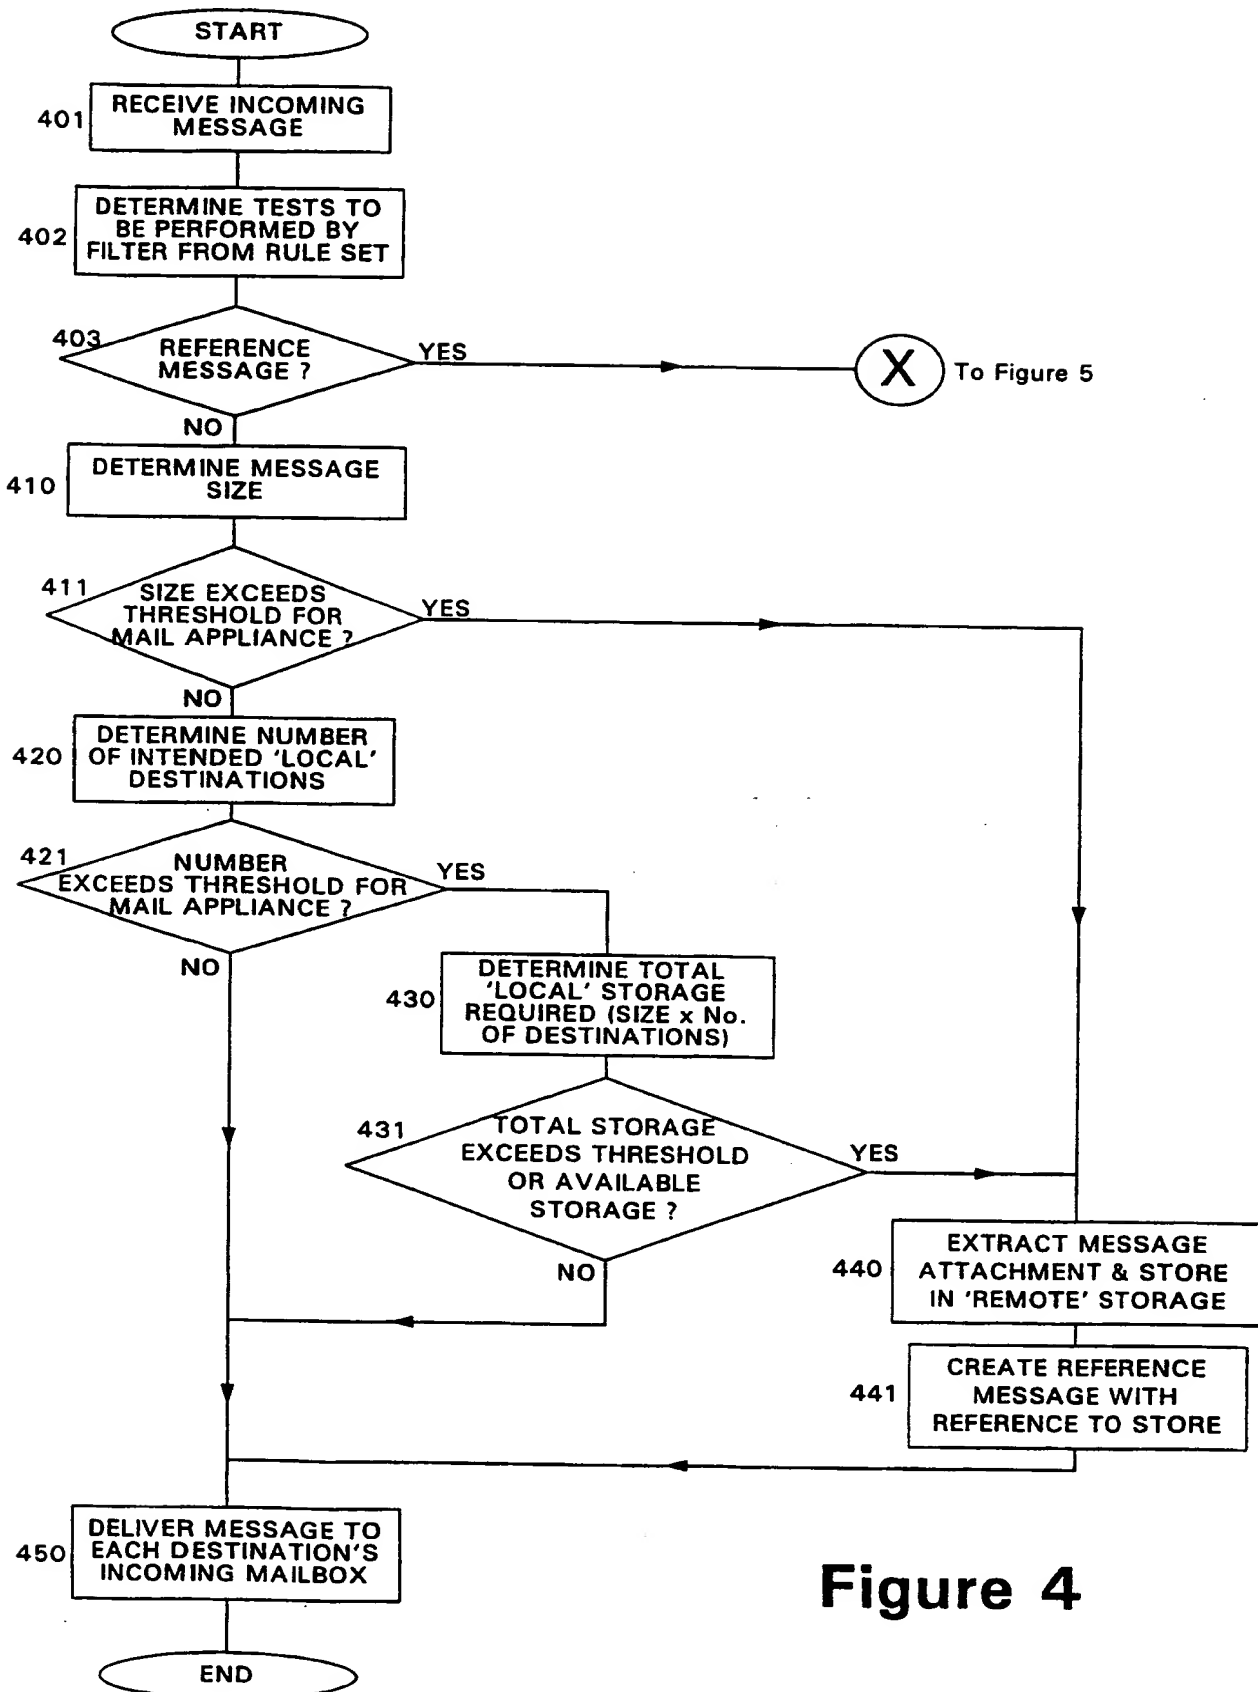
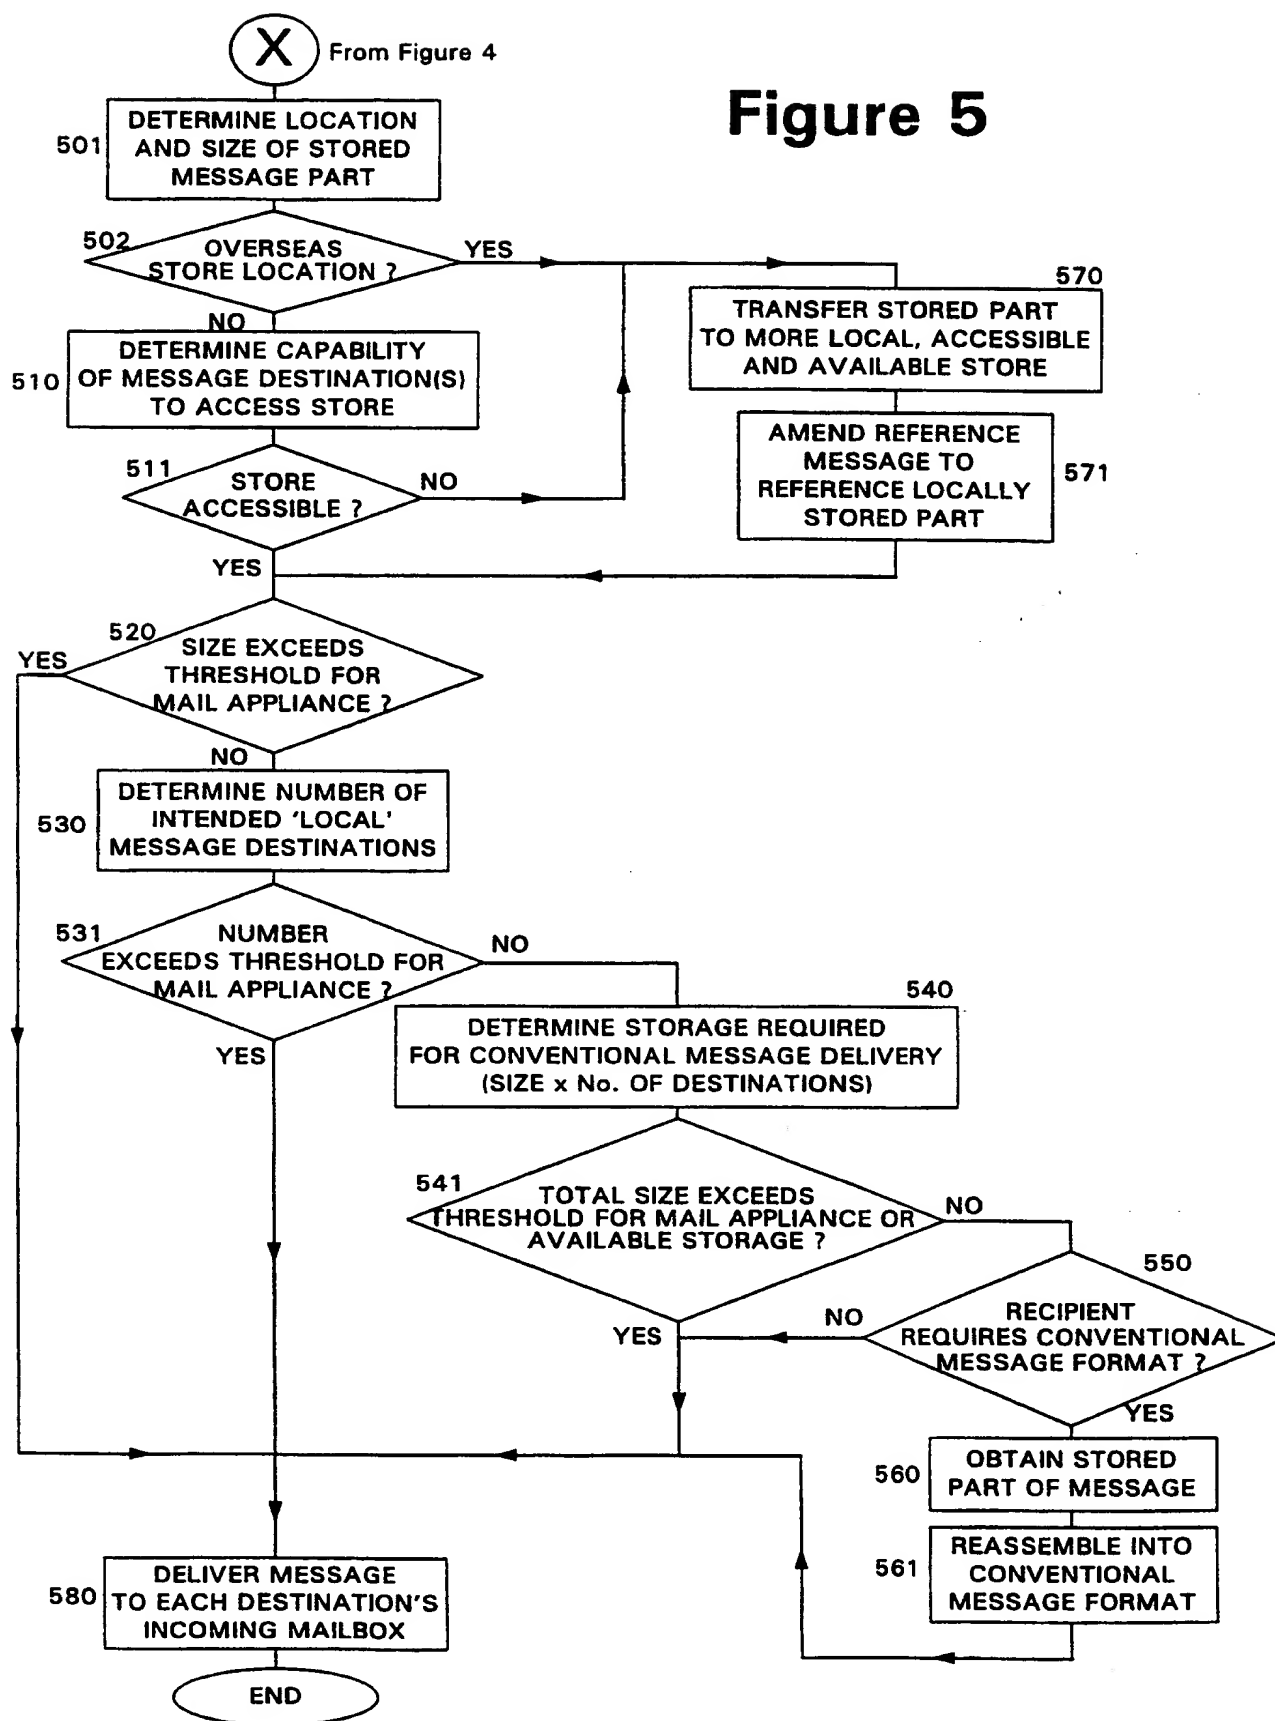


Figure 4

5/5

Figure 5



A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04L12/58 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L G06F H04M

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	BLUM C ET AL: "THE GLOBAL STORE SERVER - A MULTIMEDIA TELESERVICE COMPONENT" MULTIMEDIA TRANSPORT AND TELESERVICES. INTERNATIONAL COST 237 WORKS PROCEEDINGS, VIENNA, NOV. 13 - 15, 1994, 13 November 1994, pages 35-46, XP000585293 HUTCHISON D;DANTHINE A; LEOPOLD H; COULSON G (EDS)	1,3,5-9, 14,18
Y	see page 35, line 1 - page 42, line 15 see page 45, paragraph 6 see figures 1-3	4,10
X	US 5 548 789 A (NAKANURA YASUO) 20 August 1996 see column 7, line 42 - column 8, line 21	1-3,5,7, 8,14,18
Y	see column 7, line 19 - line 27 see figures 8,9	4,9,10
	--- -/-	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"S" document member of the same patent family

Date of the actual completion of the international search

26 November 1998

Date of mailing of the international search report

04/12/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Karavassilis, N

INTERNATIONAL SEARCH REPORT

Inter national Application No

PCT/GB 98/02696

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	GRIESMER S J ET AL: "EVOLUTION OF MESSAGING STANDARDS" AT & T TECHNICAL JOURNAL, vol. 73, no. 3, 1 May 1994, pages 21-45, XP000452252	4
A	(Especially page 37, col. 2 line 27 to page 40, col. 1, line 15, "bulk mailing lists") see the whole document	1,3,5-18
Y	BOGEN M ET AL: "W3GATE - A WEB ACCESS FOR OUTSIDERS" COMPUTER NETWORKS AND ISDN SYSTEMS, vol. 28, no. 14, November 1996, pages 1979-1990, XP000630956	10
A	see the whole document	1-9, 11-18
Y	US 5 606 693 A (NILSEN KENNETH ET AL) 25 February 1997 see column 1, line 1 - column 2, line 47; figure 1	9
A	THIMM H ET AL: "A MAIL-BASED TELESERVICE ARCHITECTURE FOR ARCHIVING AND RETRIEVING DYNAMICALLY COMPOSABLE MULTIMEDIA DOCUMENTS" MULTIMEDIA TRANSPORT AND TELESERVICES. INTERNATIONAL COST 237 WORKS PROCEEDINGS, VIENNA, NOV. 13 - 15, 1994, 13 November 1994, pages 14-34, XP000585292 HUTCHISON D;DANTHINE A; LEOPOLD H; COULSON G (EDS) see page 15, line 21 - line 27 see page 15, line 43 - page 22, line 15 see page 29, paragraph 6.2 see page 31, paragraph 7 see figures 2,4,5,7	11-13, 15-17
A	CARRIER S ET AL: "PRACTICAL MULTIMEDIA ELECTRONIC MAIL ON X.400" IEEE MULTIMEDIA, vol. 2, no. 4, 21 December 1995, pages 12-23, XP000542075	
A	"IS YOUR NETWORK READY FOR NOTES? HEADS-UP DESIGN AND CLOSE ATTENTION TO TRAFFIC FLOW CAN KEEP THE WORLD'S BEST-SELLING GROUPWARE FROM BRINGING THE ENTERPRISE TO ITS KNEES" DATA COMMUNICATIONS, vol. 24, no. 5, 1 April 1995, pages 83-86, 88, 90, 92, XP000501583	
	--- -/--	

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 98/02696

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 720 333 A (AT & T CORP) 3 July 1996 see page 2, column 2, line 5 - line 12 see page 2, column 2, line 33 - line 36 see abstract	1,11,18
A	----- "METHODOLOGY FOR MAIL FILTRATION PROBES" IBM TECHNICAL DISCLOSURE BULLETIN, vol. 34, no. 9, 1 February 1992, page 372 XP000301914 see the whole document -----	1,11,18

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 98/02696

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5548789 A	20-08-1996	JP 4280140 A	06-10-1992
		JP 5136818 A	01-06-1993
US 5606693 A	25-02-1997	US 5668986 A	16-09-1997
EP 0720333 A	03-07-1996	US 5619648 A	08-04-1997
		CA 2159973 A	31-05-1996
		JP 8263404 A	11-10-1996

THIS PAGE BLANK (USPTO)